Fermi National Accelerator Laboratory

TM-1596

# The Fermilab Central Computing Facility Architectural Model*

J. Nicholls
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois

May 1989

* Presented at the 1989 Conference on Computing in High Energy Physics, Oxford, England, April 10-14, 1989.

Operated by Universities Research Association, Inc., under contract with the United States Department of Energy

# The Fermilab Central Computing Facility Architectural Model

J. Nicholls
Fermilab

## Abstract

The goal of the current Central Computing Upgrade at Fermilab is to create a computing environment that maximizes total productivity, particularly for high energy physics analysis. The Computing Department and the Next Computer Acquisition Committee decided upon a model which includes five components: an interactive front end, a Large-Scale Scientific Computer (LSSC, a mainframe computing engine), a microprocessor farm system, a file server, and workstations. With the exception of the file server, all segments of this model are currently in production: a VAX/VMS Cluster interactive front end, an Amdahl VM computing engine, ACP farms, and (primarily) VMS workstations.

This presentation will discuss the implementation of the Fermilab Central Computing Facility Architectural Model. Implications for Code Management in such a heterogeneous environment, including issues such as modularity and centrality, will be considered. Special emphasis will be placed on connectivity and communications between the front-end, LSSC, and workstations, as practiced at Fermilab.

## 1. The Fermilab Model

The following is from the Report of the Fermilab Next Computer Acquisition Committee dated December 18, 1984.

> The goal of the FY86/87 acquisition is to create a computing environment that maximizes total productivity, particularly for high energy physics programs. The system requires large CPU power to run production programs. It also requires user friendly tools for program development (including e.g., editors, library tools, etc.), provisions for communications with other computers (both on site and off site) and modern interactive tools to aid in the various computing activities.

The committee concluded that there was no single affordable solution, and the committee's recommendations led to an architecture with five major components:

- Front-end user interface system
- Large mainframe computing engine
- Farm-like system of parallel processors or attached processors
- File Server
- Workstations

The following diagram illustrates the model.

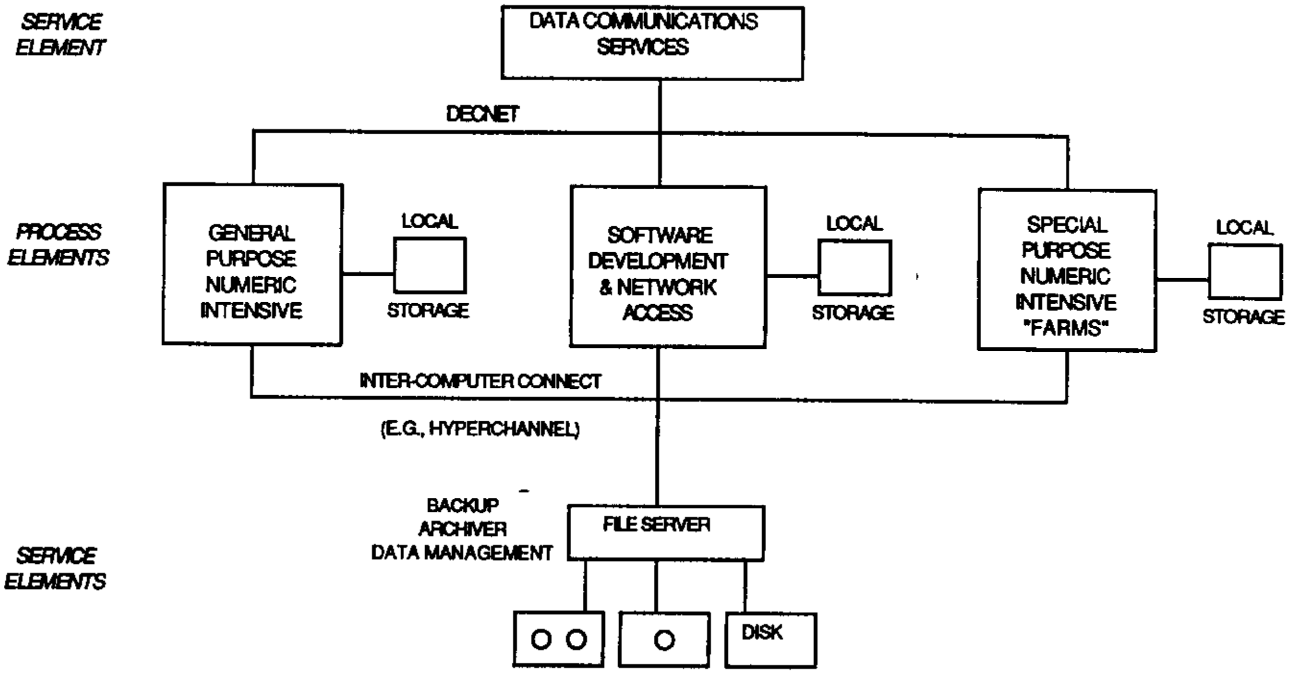


**Figure 1. Fermilab Computing Systems Architecture**

This model allows the laboratory the flexiblility to pursue new architectures and opportunities with relative ease and speed. It also has implications for the computer users, some of which are discussed below.

## 2. Implementation of the Model

*You've gotta develop your solution to your problem.*
*John Manzo, DEC, Asilomar*

As part of the Central Computing Upgrade Project, the laboratory has acquired, at this point, an Amdahl 5890/600, additional DEC VAX equipment of the 8600/8800 class, and additional ACP equipment. The file server piece will be purchased by competitive bid in this fiscal year. Workstations distributed throughout the laboratory have been purchased out of individual budgets. The diagram below illustrates the current computing configuration.

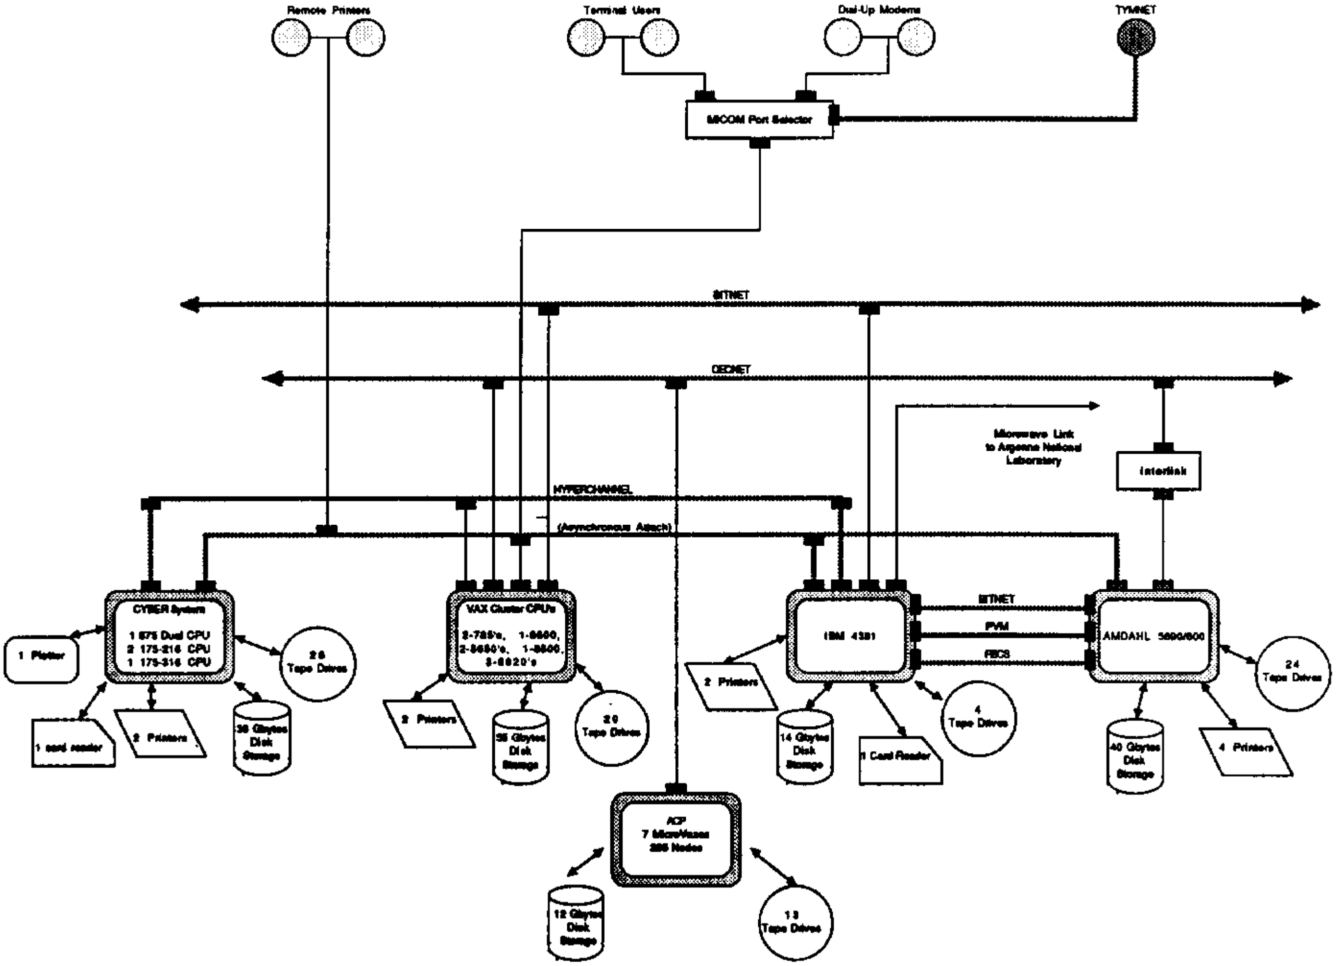


**Figure 2. Central Facility Computing Diagram**

## <u>Introduction of the Amdahl</u>

The integration of the Amdahl into the Fermilab environment has been one of the major challenges of the department in the last year. Just as the model was developed by an official committee of users, the introduction of the Amdahl was carried out in close cooperation with a small group of users (called "early birds").

The department chose to run an un-modified VM system, which cut us off from easy access to a large part of the software tools available in the high energy physics community for VM. Furthermore, the model's role for the VM system was not as a general-purpose system, but as a computing engine, and the model, as well as the community, required that the Amdahl have excellent connectivity to the VAX VMS world.

The following list includes some of the important actions we took over the last year to fit the Amdahl into our environment:

- Installed products to make use from the VAX easy and comfortable
- Installed Interlink for DECnet connection to the VAX world
- Obtained/converted HEPVM products such as GIME, INSTALL, MAIL
- Wrote a comprehensive Amdahl User's Guide (over 300 pages)
- Tailored classes to our environment and high energy physics needs
- Installed VM Software Products for batch, tape, security, backup, archive, etc.

## 3. The Current Environment

*The average user can remember only 1.1 operating systems.*
*CERN user meeting in 1970*

The evidence indicates that the model was suitable for our environment, and that it has, on the whole, been accepted.

### VAX Clusters

Software development for VAX, Amdahl, and ACP primarily takes place on the VAX cluster. It continues to be a very popular system, to the point of being "loved to death." VMS is considered easy to learn and VMS expertise is widespread. It is the lab general-purpose resource for software development and verification, electronic mail and conferencing, word processing, small batch jobs, and many other uses. Tools are available on the clusters for submitting and monitoring batch jobs which are running on the Amdahl and ACPs. There are more than 400 users logged in to the two central clusters on a "typical" afternoon.

The tools for program development and debugging are excellent. A very popular code management combination at Fermilab is CDF's EXPAND and DEC's CMS.

One of the major implications of the model for the users relates to code development. Programmers must adhere to a subset of FORTRAN that is widely transportable. Several experiments have had very good experiences porting their code to the Amdahl and the ACP because they had rather strict coding rules. Other experiments have had more difficulty, and some Cyber applications remain to be converted.

### Amdahl

The Amdahl 5890/300 was made available to the users in September, and was 100% utilized before it had been in place for 6 months. On the other hand, interactive usage of it is low (30 to 40 users in the daytime), which is consistent with the Fermilab model. Two of the major experiments using the Amdahl use it almost exclusively from the VAX cluster. None of the experiments using the Amdahl recommend widespread expert knowledge of VM in their groups.

### ACP I Farms

The current total of ACP I farms includes 7 microVAXes, over 400 nodes, 12 Gigabytes of disk storage, and 13 tape drives, with a total power of 300 VAX 780 equivalents. They are 100% utilized, with 7 experiments as major users, and several more as minor users. There are also nodes connected to the clusters for program development.

### Workstations

The cluster response time is one of the reasons why workstations are now a very popular purchase at Fermilab, but not the only one. There are now more than 200 DEC workstations, currently the predominant variety, at the laboratory.

### Networks

Interlink's IBMvm/DECnet gateway makes the Amdahl a node on DECnet. The quality of the networking between the systems is a major key to the success of the Fermilab model, and in particular the transparency, flexibility, and reliability of the Interlink connection has made it possible for the Amdahl to be effectively used remotely from the VAXes.

Network Systems HYPERchannel connects the Cyber with the other computers and acts as an alternate path for transfer between the VAX clusters and the Amdahl.

BITNET is used for communication with the rest of the world, as well as an alternate path between the VAX clusters and the Amdahl.

### Other Resources

There are more than 50 Talaris laser printers at the laboratory, most of them easily accessible from the central clusters, the Amdahl, the Cybers, and workstations.

## 4. The Future

> *What is really happening is that computing is becoming a network wire with resources hanging off of it.*
>
> *David Ritchie, Fermilab*

We believe one of the characteristics of the model is it allows us to be poised to take advantage of new opportunities. The following is a list of some of the current and near-term projects.

### Central Facility

We are currently implementing VM/XA to support the 4 CPU's of the 5890/600, and later to exploit access to more than 16 Megabytes of memory.

Cyber decommissioning will occur at the end of fiscal year 1990, and a major upcoming project is to convert the remaining Cyber users.

## Distributed Computing

The Computing Department provides considerable support for workstations and local area VAX clusters including: networking and management planning advice, systems support, software distribution support and tools, licensing and management support. This support is directed at 22 clusters already. These clusters are dedicated to varied uses ranging from general physics analysis to engineering to data acquisition.

## Parallel Processing

Event parallelism is a proven technology at Fermilab, providing cost-effective computing for tasks such as track reconstruction. The lab is now experimenting with new parallel architectures such as UNIX farms based on the MIPS chip, VMS farms, both in-house versions and commercial solutions.

## Common File Server

The last piece of the upgrade is a common file server. The common file server is both an important tool in supporting the heterogenous computing environment, as well as an important component in creating a "transparent" computing environment as seen by the user. The common file server is attached to each of the process elements (Figure 1) and also serves to host the growing variety of storage media.

## New Media Opportunities

The Central Facility currently provides an Exabytes 8 mm tape copy facility, primarily for providing cheap easy-to-use media for use at home institutions.

IBM 3480 cartridges are supported on the Amdahl and may become popular, especially when the density is increased.

## Software Support

There are a number of areas where lack of money and manpower have limited our effectiveness in providing and supporting productivity tools. In addition, the current data base systems available are not adequate for all of the tasks. And finally, distributed computing provides altogether new requirements for a centrally supported program library system and other services.